

(19)



Europäisches Patentamt

European Patent Office

Office européen des brevets

(11) Publication number:

0 327 203
A2

(12)

EUROPEAN PATENT APPLICATION

(21) Application number: 89300139.6

(51) Int. Cl.4: G06F 13/14 , G06F 13/16

(22) Date of filing: 09.01.89

(30) Priority: 05.02.88 US 152779


(43) Date of publication of application:
09.08.89 Bulletin 89/32(84) Designated Contracting States:
DE FR GB(71) Applicant: International Business Machines
Corporation
Old Orchard Road
Armonk, N.Y. 10504(US)(72) Inventor: Rodiger, William King,
18L Sunset Gardens Apartments
Kingston New York 12401(US)
Inventor: Thorson, Jon Edwin
Kelly Circle
Poughkeepsie New York 12601(US)
Inventor: Ting, Yee-Ming
31 Ronduit Harbour
Port Ewen New York 12466(US)(74) Representative: Grant, Iain Murray
IBM United Kingdom Limited Intellectual
Property Department Hursley Park
Winchester Hampshire SO21 2JN(GB)

(54) NxM arbitrating non-blocking high bandwidth switch.

(57) An N x M matrix adapted to couple N inputs from N processor to M basic storage modules is disclosed. The system includes arbitrators and gating means for each output responsive to request signals for simultaneously coupling data from a plurality of processors to requested basic storage modules under arbitrator control.

EP 0 327 203 A2

BEST AVAILABLE COPY



NxM ARBITRATING NON-BLOCKING HIGH BANDWIDTH SWITCH

This invention relates to NxM arbitrating non-blocking high bandwidth switches for use, for example, in multiprocessing systems as arbitration systems for handling requests for access from a plurality of inputs to a plurality of outputs.

It is desirable in large complex computer systems to have multiple memories or basic storage modules (BSM) and have multiple processors communicating with multiple memories with buffering.

In the prior art, there is known arbitrating systems such as is disclosed in US-A-4 473 880 of Budde, et al or US-A-4 499 538 of Finger, et al that provide some form of arbitration system to several processors or microprocessors with a common bus. These arbitration systems with a common bus are relatively slow systems in that only an input to the bus from one unit can be applied via the bus to an output unit in a given time cycle. Cross point switches such as in a switching matrix as described in US-A-4 417 245 of Melas, et al couple multiple inputs to multiple outputs simultaneously provided a given input does not want to conflict with another given input at the same output and assumes that some form of separate control is provided.

There is no arbitrator involved.

In seeking to accommodate both uncontrolled input requests and arbitration, the present invention provides a NxM arbitrating non-blocking high bandwidth switch comprising:

M output ports;

N data inputs and request inputs, each request input being associated with a data input and indicating to which of the M output ports the associated data is to be applied;

M arbitrators and gating means, each arbitrator and gating means being coupled between the N data inputs and request inputs and a different one of the M output ports; and

each of the arbitrator and gating means being responsive to the data inputs and request inputs to gate data inputs simultaneously to the requested output ports when there is no contending request inputs for the same output port but, when there are contending request inputs, to sequence the associated contending data inputs to the associated output port or ports, while simultaneously applying the remaining data inputs directly.

The present invention equally provides a multiprocessor system including such a switch.

In one embodiment of the present invention, disclosed hereinafter, a plurality of input signals comprising data and associated output request code are selectively applied to one of a plurality of

ling each of the outputs. The arbitrators and gating means are responsive to the input signals for gating the non-contending input signals simultaneously to the non-contending outputs and for sequencing any contending signals to the associated outputs.

The present invention will be described further by way of example with reference to an embodiment thereof as illustrated in the accompanying drawings, in which:

Figure 1 is a block diagram of an overall system according to one embodiment of the present invention;

Figure 2 is a block diagram of an arbitrator thereof; and

Figure 3 is a block diagram of a preferred manner of extending the system of Figure 1 for additional processors and/or memories.

Referring to the system block diagram of Figure 1, there is illustrated one example of an N x M arbitrating non-blocking high bandwidth switch in a multiprocessor system. The system shown in Figure 1 refers to processors 11 through 14 being selectively coupled to one of four basic storage modules (BSM) 21 through 24. In actual practice, the number of processors connected to basic storage may be 32 or greater. The inputs from the processors 11 through 14 are applied to respective buffers 31 through 34. These buffers 31 through 34 may be for example FIFO (First In, First Out) buffers. The output from each of the processors 11 through 14 includes a request code signal and associated data signals or data packet. The request code signal could be simply a logic 1 level at one of the four lead wires. The output request code signal from each of the buffers 31 through 34 is applied to each of four arbitrators 41 through 44. The request code signal and associated data signals or data packet are sequenced together through the buffers 31 through 34. The request code signal could be via four wires or lines, with only one of the four wires being at a logic 1 level and the other three wires being at a logic zero level. The logic 1 level may be five volts where the logic zero may be at ground potential. If, for example, processor 11 wants to be coupled to BSM 21, the processor 11 provides a logic 1 level on lead 101 from buffer 31 and logic 0 level to leads 102 through 104 at the first request ports of arbitrators 42 through 44. If processor 11 wants to be coupled to BSM 22, BSM 23, or BSM 24, the logic 1 level would only be on line 102, 103, or 104 respectively.

BEST AVAILABLE COPY

coupled to BSM 21, BSM 22, BSM 23, or BSM 24, the logic 1 level would only be on line 141, 142, 143, or 144 respectively. Likewise, processors 12 and 13 can be coupled to any one of the basic BSM modules 21 through 24 by appropriate logic 1 level at the request line to arbitrators 41 through 44 respectively.

Each arbitrator first determines if there is an input from any of the four FIFO buffers 31 through 34. If there is more than one input to that arbitrator, it sequences the data inputs each clock cycle through an associated selector 51 through 54. The data is sent to all four selectors 51 through 54. If there are no contentions or more than one processor trying to use a given storage module at the arbitrator, the input from each buffer 31 through 34 at each selector 51 through 54 is simultaneously coupled to the associated BSM 21 through 24 via associated buffer 61 through 64. Selectors 51 through 54 are associated with each of the BSM 21 through 24. The arbitrator 41 through 44 identify which output BSM 21 through 24 the input data from a given processor is to be applied and provide a select code to the selectors 51 through 54 to gate the data from buffers 31 through 34 to the appropriate one of buffers 51 through 64.

If there is more than one contender for a given BSM, one of the arbitrators 41 through 44 associated with a given BSM determines the sequence out of the associated selectors 51 through 54. For example, if there is a request code signal from processor 11 and from processor 12 via buffers 31 and 32 for BSM 21, they are both applied to the arbitrator 41. The arbitrator 41 provides, for example, a first select code to the selector 51 during the first clock pulse to gate the output from the FIFO buffer 31 through the selector 51 to buffer 61 and to the BSM 21. During the next clock cycle, the second select signal code is provided to selector 51 and processor 12 data at buffer 32 is coupled through the selector 51 to the BSM Module 21.

In Figure 2, there is illustrated a logic block diagram for each arbitrator of Figure 1. There are four input request ports 201 through 204 for each arbitrator 41 through 44. The first input port 201 is coupled to processor 11 via buffer 31. For arbitrator 41, this input port 201 is coupled via wire lead 101. The second input port 202 is coupled to processor 12 via buffer 32, the third port 203 is coupled to processor 13 via buffer 33 and the fourth port 204 is coupled to processor 14 via buffer 34. The request at inputs 201 through 204 are applied to the respective AND gates A1 through A4. The outputs from AND gates A1 through A4 are applied to the set inputs of flip-flop registers S1 through S4. The Q output from register S1 is applied directly to the output lead 205 and to NOR gate 100. The Q outputs from the registers

S2 through S4 are also applied to the NOR gate 100. The \overline{Q} output of switching register S1 is applied to AND Gate A5. The Q output from register S2 is also applied to the AND Gate A5. The Q output of switches S3 and S4 are coupled to the respective inputs of AND Gates A6 and A7. AND Gate A6 also receives the \overline{Q} output from registers S1 and S2. AND Gate A7 receives a \overline{Q} output from registers S1, S2 and S3 and the Q output from register S4. The Q output from switch S1 is applied to the reset input of switch S1. The output of AND Gate A5 is coupled to the reset input of register S2. The output of AND Gate A6 coupled to the reset input of register S3 and the output of AND Gate A7 is coupled to the reset input of register S4. The outputs of AND Gates A5 through A7 on leads 206 through 208 and the Q output of register S1 on lead 205 provide a four-bit address code which is applied to the corresponding selector 51 through 54 to select the output. In actual practice for the simple embodiment of Figure 1 with only four processors and four BSM modules that one of the arbitrator outputs 205 through 208 having a logic 1 level would enable the associated processor data to the associated basic storage module. For example, if the output from register S1 is at the logic 1 level, the processor 11 output will be coupled to the basic storage module associated with the arbitrator. If at arbitrator 41, processor 11 data would be coupled to BSM 21. If at arbitrator 42 processor 11 data is coupled to BSM 22, if at arbitrator 43 processor 11 data is coupled to BSM 23, and if at arbitrator 44 processor 11 data is coupled to BSM 24. If at arbitrator 41 there is a logic 1 level at output 208 of arbitrator 41, the data from processor 14 would be coupled to the basic storage module 21. If the logic 1 level is at output 208 of arbitrator 42, the data from processor 14 is coupled to BSM 22. If at arbitrator 43 from processor 14 to BSM 23, etc.

In the start-up state of the arbitrator, there are all logic zeros at the Q outputs and logic 1 levels at the \overline{Q} outputs of registers S1 through S4. The NOR Gate 100 provides a logic 1 level to AND Gates A1 through A4 with all zeros at its input. If there is only a request or logic 1 level at input 202, for example, this output enables AND Gate A2 with a logic 1 coupled to the set input of switch S2 to thereby provide a logic 1 level at the Q output of register S2. The logic 1 level output from register S2 is coupled to the NOR Gate 100 and in response to this logic 1 a logic low or zero level from NOR Gate 100 is provided to AND Gates A1 through A2, stopping all further requests. The logic 1 level input at AND Gate A5 from register S2 enables the logic 1 level input from the \overline{Q} output of register S1 to provide logic 1 level at the output 206. This code with a logic 1 level only on lead 206 requests the

BEST AVAILABLE COPY

selector 51 through 54 associated with the arbitrator to couple the data from the second processor 12 to the basic storage module associated with the selector. For example, with arbitrator 41 the data from processor 12 is coupled to BSM 21.

A contention exists when there are two or more requests for a given BSM. For example, in the analysis, consider requests at inputs 201 and 203 from processors 11 and 13 for BSM 21. The arbitrator 41 sees the presence of a logic 1 level at AND Gates A1 and A3 to enable the logic 1 from NOR Gate 100 to produce a set of logic 1 levels at the set inputs of registers S1 and S3. This produces a logic 1 level at the Q outputs of registers S1 and S3. The presence of the logic 1 level at either of these "Q" outputs at the NOR Gate 100 provides a logic "0" at the AND Gates A1 through A4 which stops any further requests. A logic zero or low level is provided from the \overline{Q} outputs of registers S1 and S3 to the AND Gates A5 through A7. AND Gate A6 output remains at a logic zero with the \overline{Q} output from register S1 being low. In effect, only the request at input 201 for BSM 21 is acknowledged at the first clock cycle. As soon as the select is made to couple the data from processor 11, the register S1 is reset via lead 209 providing a logic one at the \overline{Q} output of register S1 to AND gate A6 and allowing the request at input 203 to provide a logic 1 from AND gate A6. This logic 1 at output 207 of AND gate A6 selects the data then from processor 13 to the associated BSM on the next clock cycle. After the requester at input 203 selects the data, register S3 is reset by the feedback path 210 from AND gate A6 and the NOR output from NOR Gate 100 goes high. When the NOR Gate output goes high, all of the registers S1 through S4 are off and any combination of input request can be acknowledged in the next full cycle.

In the example given in Figures 1 and 2, there were two clock cycles required to handle two simultaneous requests for the given basic storage module. Since port requests are generated every cycle, the storage of the input buffers should equal the number of processors. For example, in Figure 1 this would require four buffer stages to cover the case where all four processors would be attempting to communicate with a given basic storage module. They would be stored and sequenced over four clock cycles. In the example given, there were only four inputs processor coupling to four basic storage modules, but in a preferred application there would be, for example, 32 processors communicating with 32 or more storage modules.

The arbitrators as shown in Figure 2 always acknowledge multiple requests in round-robin order of S1 through S4 until all are satisfied; that is, in instant case the request at input 201 has the highest priority with request at input 202 second, input

at 203 third and input at 204 having the lowest priority within that arbitration cycle.

It can be seen that during one clock cycle four port requests can be distributed across the four BSMs, a maximum of four BSM selects can be serviced per clock cycle period. The minimum would be one BSM select per cycle where all port requests are for the same BSM. To handle coupling of more processors to more basic storage modules such as 32 processors to 32 BSM's, for example, is shown in Figure 3. In this arrangement, each processor 301-332 provides a 5-bit encoded request signal for each data block or packet to be applied to the BSM. The request code signal with its associated data packet is sequenced through the associated buffer. A decoder 301a-332a is at the output of each input buffer 301b-332b for decoding each 5-bit processor encoded request signal and providing a logic 1 request via lead on one of the 32 output lines to the appropriate arbitrator. For example, decoder 301a decodes the encoded output from buffer 301b indicating it's data to be applied to BSM 302f and provides a logic 1 on lead 400 or request to arbitrator 302c. The data at buffer 301b is then enabled through select switch 302d to memory BSM 302f via buffer 302e. The arbitrators 301c-332c would be like Figure 2 with 32 inputs and outputs instead of four input and outputs.

Claims

1. A NxM arbitrating non-blocking high bandwidth switch comprising:

M output ports;

N data inputs and request inputs, each request input being associated with a data input and indicating to which of the M output ports the associated data is to be applied;

M arbitrators and gating means, each arbitrator and gating means being coupled between the N data inputs and request inputs and a different one of the M output ports; and

each of the arbitrator and gating means being responsive to the data inputs and request inputs to gate data inputs simultaneously to the requested output ports when there is no contending request inputs for the same output port but, when there are contending request inputs, to sequence the associated contending data inputs to the associated output port or ports, while simultaneously applying the remaining data inputs directly.

2. A switch as claimed in claim 1, wherein the M arbitrators and gating means include arbitrators and gating means which have a priority sequence for gating contending request signals out of the gating means.

3. A switch as claimed in either preceding claim, wherein the arbitrator means produces a coded signal to the gating means which identifies the request input and the gating means is coupled directly to the data inputs and in response to the code signal from the arbitrator gates that input to the associated output port.

4. A multi-processor system for randomly coupling an N plurality of processors to an M plurality of basic storage modules (BSM) via a switch as claimed in any preceding claim, wherein each basic storage module is connected to its own output port; each processor is connected to its own input port and provides thereto, as required, a data packet and a request code signal identifying which basic storage module the data packet is to be applied;


whereby, when there is only one data packet targetted at a particular basic storage module in any cycle, and thus, there is only one request applied to the corresponding arbitrator, the arbitrator causes the associated gating means to gate the data package directly to its coupled basic storage module and when more than one request is received at a given cycle targetted at the same basic storage module, and, hence, more than one request is received at the corresponding arbitrator, such arbitrator causes sequencing of the counterpart data packets to the associated basic storage module.

5. A system as claimed in claim 4, including a buffer means coupled between the processors and the arbitrators and gating means for storing and sequencing the data packets and request signals each clock cycle.

6. A system as claimed in claim 5, wherein the first buffer means has at least as many stages as there are arbitrators.

7. A system as claimed in claim 5 or claim 6, including a second buffer means between each gating means and associated Basic Storage Module.

BEST AVAILABLE COPY



BEST AVAILABLE COPY

BE




FIG. 3A

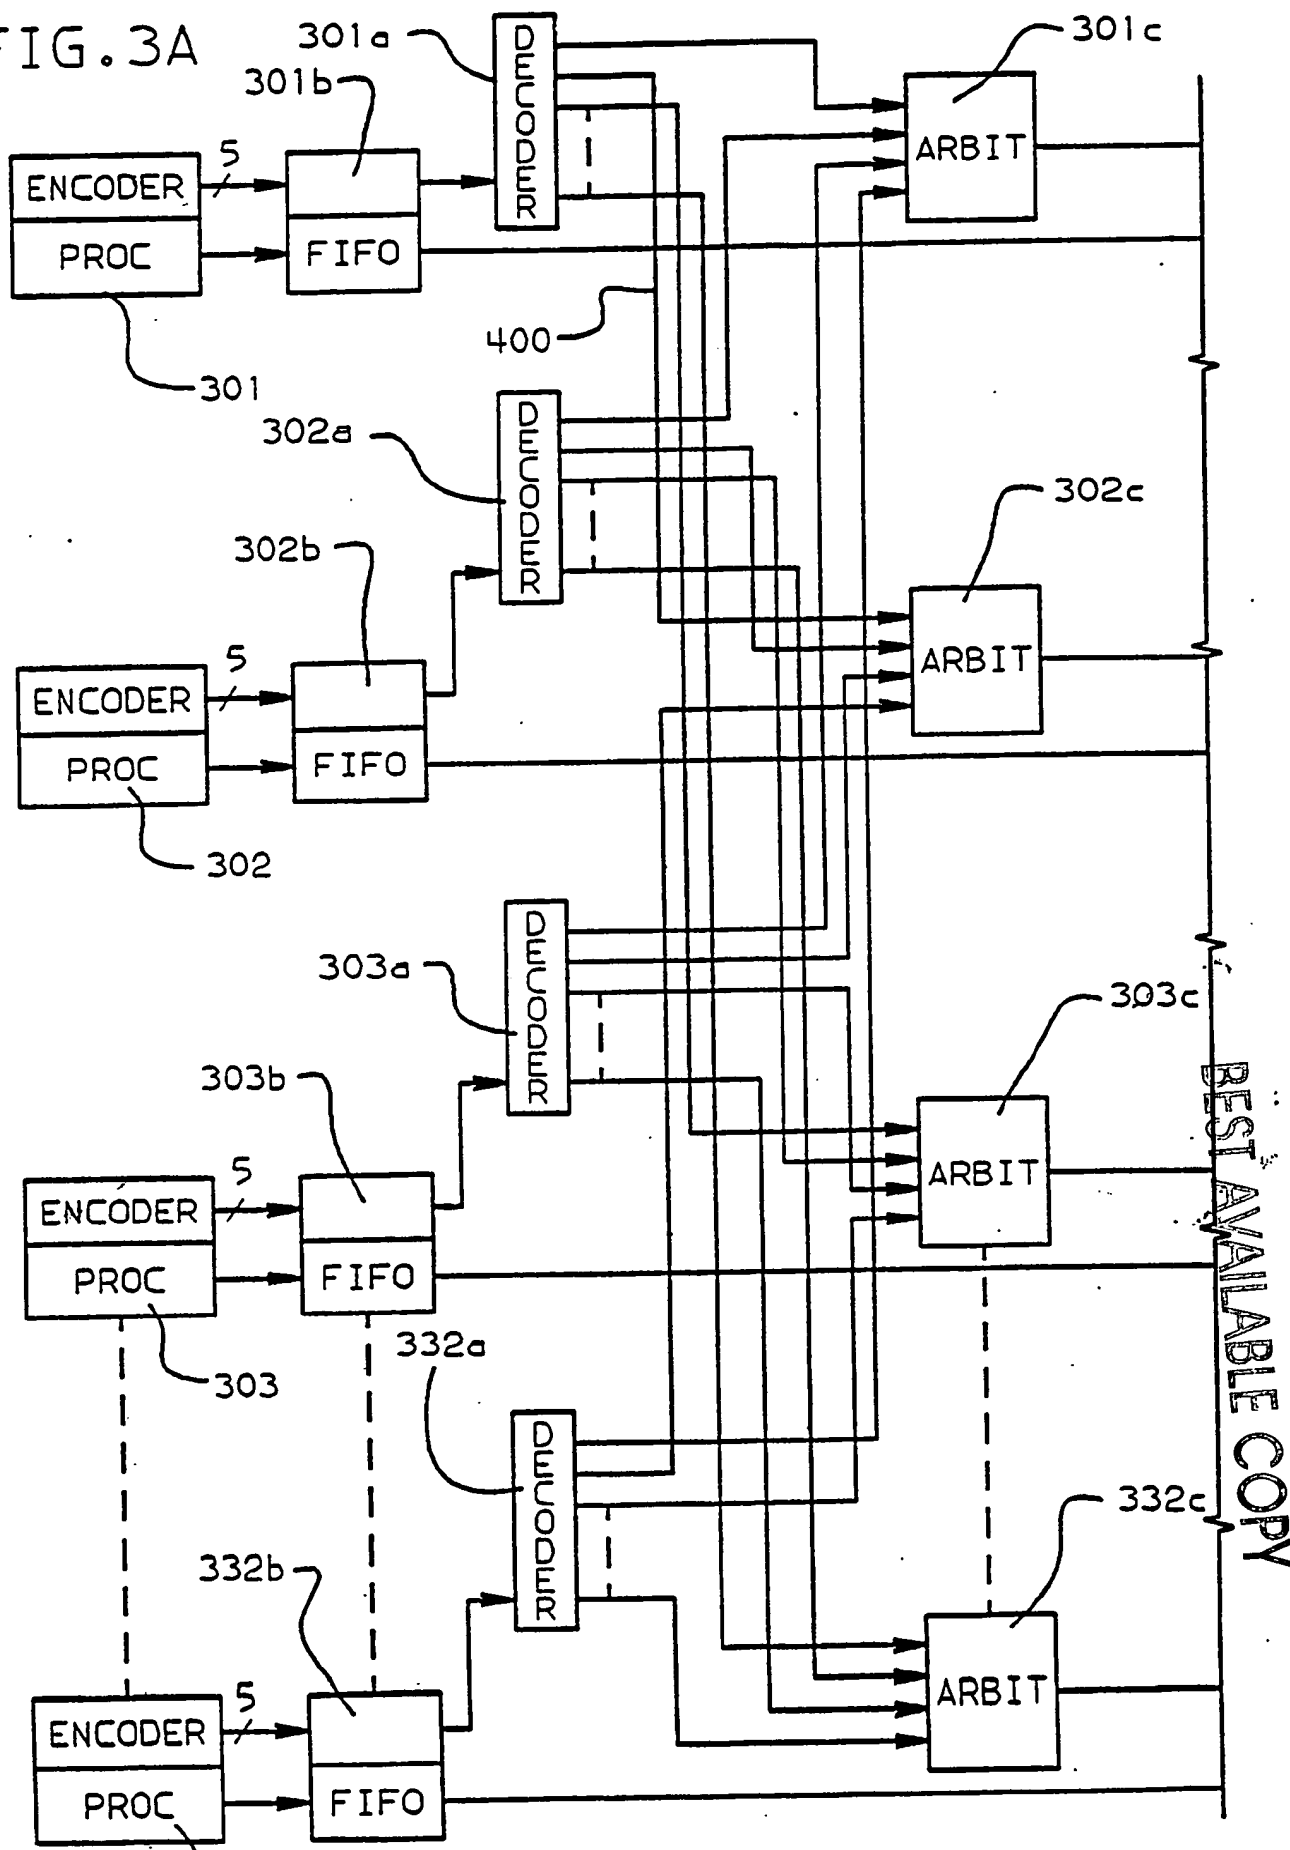



FIG. 3B



BEST AVAILABLE COPY

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) Publication number:

0 327 203 A3

(12)

EUROPEAN PATENT APPLICATION

(21) Application number: 89300139.6

(51) Int. Cl.⁵: G06F 13/14, G06F 13/16

(22) Date of filing: 09.01.89

(30) Priority: 05.02.88 US 152779

(43) Date of publication of application:
09.08.89 Bulletin 89/32(64) Designated Contracting States:
DE FR GB(68) Date of deferred publication of the search report:
02.05.91 Bulletin 91/18(71) Applicant: International Business Machines
Corporation
Old Orchard Road
Armonk, N.Y. 10504(US)


(72) Inventor: Rodiger, William King
18L Sunset Gardens Apartments
Kingston New York 12401(US)
Inventor: Thorson, Jon Edwin
Kelly Circle
Poughkeepsie New York 12601(US)
Inventor: Ting, Yee-Ming
31 Rondult Harbour
Port Ewen New York 12466(US)

(74) Representative: Atchley, Martin John
Waldegrave
IBM United Kingdom Limited Intellectual
Property Department Hursley Park
Winchester Hampshire SO21 2JN(GB)

(54) NxM arbitrating non-blocking high bandwidth switch.

(57) An N x M matrix adapted to couple N inputs from N processor to M basic storage modules is disclosed. The system includes arbitrators and gating means for each output responsive to request

signals for simultaneously coupling data from a plurality of processors to requested basic storage modules under arbitrator control.



EP 0 327 203 A3

BEST AVAILABLE COPY



European
Patent Office

EUROPEAN SEARCH REPORT

Application Number

EP 89 30 0139

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int. Cl.5)
X,Y	EP-A-0 097 351 (NEC CORPORATION) * page 5, line 13 - page 6, line 18 * * page 10, line 19 - page 12, line 15; figures 1, 6 * - - - -	1-3,4,5,7	G 06 F 13/14 G 06 F 13/16
Y	IEEE TRANSACTIONS ON COMPUTERS. vol. C-34, no. 10, October 1985, NEW YORK US pages 918 - 926; C.R.DAS et al: "Bandwidth Availability of Multiple-Bus Multiprocessors" * page 918, right-hand column, lines 14 - 46; figures 1-3 * - - - -	4,5,7	
A	GB-A-1 170 434 (GENERAL ELECTRIC COMPANY) * page 2, lines 18 - 113; figure * - - - -	1,4	
A	US-A-3 096 447 (J.E.HILL ET AL) * column 2, line 22 - column 3, line 4; figure 1 * - - - - -	1,4	
			TECHNICAL FIELDS SEARCHED (Int. Cl.5)
			G 06 F
The present search report has been drawn up for all claims			
Place of search		Date of completion of search	Examiner
The Hague		05 March 91	WANZEELE R.J.
CATEGORY OF CITED DOCUMENTS X: particularly relevant if taken alone Y: particularly relevant if combined with another document of the same category A: technological background O: non-written disclosure P: intermediate document E: earlier patent document, but published on, or after the filing date D: document cited in the application L: document cited for other reasons &: member of the same patent family, corresponding document			